

Audio-Visual Floorplan Reconstruction

Senthil Purushwalkam*¹ Sebastia Vicenc Amengual Gari² Vamsi Krishna Ithapu²
Carl Schissler² Philip Robinson² Abhinav Gupta³ Kristen Grauman^{3,4}

¹Carnegie Mellon University ²Facebook Reality Labs ³Facebook AI Research ⁴University of Texas at Austin

Abstract

Given only a few glimpses of an environment, how much can we infer about its entire floorplan? Existing methods can map only what is visible or immediately apparent from context, and thus require substantial movements through a space to fully map it. We explore how both audio and visual sensing together can provide rapid floorplan reconstruction from limited viewpoints. Audio not only helps sense geometry outside the camera’s field of view, but it also reveals the existence of distant freespace (e.g., a dog barking in another room) and suggests the presence of rooms not visible to the camera (e.g., a dishwasher humming in what must be the kitchen to the left). We introduce AV-Map, a novel multi-modal encoder-decoder framework that reasons jointly about audio and vision to reconstruct a floorplan from a short input video sequence. We train our model to predict both the interior structure of the environment and the associated rooms’ semantic labels. Our results on 85 large real-world environments show the impact: with just a few glimpses spanning 26% of an area, we can estimate the whole area with 66% accuracy—substantially better than the state of the art approach for extrapolating visual maps.

1. Introduction

Floorplans of complex 3D environments—such as homes, offices, shops, churches—are a compact ground-plane representation of their overall layout, showing the different rooms and their connectivity. Floorplans are useful for visualizing a large space, navigating an unfamiliar building, planning safety routes, and communicating architectural designs.

Traditionally a floorplan is created by distilling a *fully observed* 3D environment into its footprint—whether manually or with the aid of 3D sensors [42, 31]. Recent research aims to infer room layouts using imagery and/or scans, with impressive results [29, 8, 27, 46]. However, existing methods are limited to mapping the regions they directly observe.

*work done while interning at Facebook AI Research. Project webpage: <http://www.cs.cmu.edu/~spurushw/publication/avmap>




Figure 1: Audio-visual floorplan reconstruction: A short video walk through the house can reconstruct the visible portions of the floorplan, but is blind to many areas. We introduce *audio-visual floorplan reconstruction*, where sounds in the environment help infer both the geometric properties of the hidden areas as well as the semantic labels of the unobserved rooms (e.g., cooking sounds behind a wall to the camera’s left suggest the kitchen).

They require a dense walk-through for the camera to capture most of the space. This is often wasteful if not impossible for a human manually photographing a large space or a robot needing to immediately perform tasks in a new environment. In this work, we address this shortcoming by estimating floorplans beyond the regions captured by a camera.

Our idea is to “see” beyond the visible regions by listening. Audio provides strong spatial and semantic signals that complement the mapping capabilities of visual sensing. In particular, the value of audio for floorplan estimation is threefold. First, observed sound is inherently driven by geometry; audio reflections bounce off major surfaces and reveal the shape of a room, beyond the camera’s field of view. Second, sounds heard from afar can suggest the existence of distant freespace where the sounding object could exist. Third, hearing semantically meaningful sounds from different directions naturally reveals the plausible room layouts based on the activities or objects those sounds represent.

For example, a shower running suggests the direction of the bathroom, even before we see it; microwave beeps suggest a kitchen; climbing footsteps suggest a staircase. See Figure 1.

To this end, we propose a new research direction: *audio-visual floorplan reconstruction*. Given a short RGB video complete with multi-channel audio, the goal is to produce a 2D floorplan that shows the freespace and occupied regions and divides them into a discrete set of semantic room labels (family room, kitchen, etc.). Importantly, the floorplan output extends significantly beyond the area directly observable in the video frames. This efficiency is critical for navigating robots that need to act without exhaustively touring a space, as well as offline scenarios where a user wants to extract a broad map from an existing short video. For example, imagine floorplans inferred from YouTube videos to facilitate real estate viewing or from wearable camera video to benefit spatial reasoning in augmented reality.

Our AV-Map approach works as follows. We devise a deep convolutional neural network architecture that leverages sequences of audio and visual data to reason about the structure and semantics of the floorplan. Our model independently extracts floorplan-aligned features from audio and RGB data, encodes sequences of features of each modality using self-attention mechanisms, and finally fuses information from audio and RGB using a decoder architecture.

We consider two settings: device-generated sounds (active) and environment-generated sounds (passive). In the active setting, the camera emits a known sound while it moves. This corresponds to a use case where a person or robot does a swift walk-through of an environment while their phone/camera emits some sound. In the passive setting, we observe only naturally occurring sounds made by objects and people in the building. This corresponds to a use case where we are simply given a passively recorded video, likely captured for some other purpose.

To our knowledge, ours is the first attempt to infer floorplans from audio-visual data. Our results on 85 large real-world, multi-room environments show that AV-Map not only consistently outperforms traditional vision-based mapping, but also improves the state-of-the-art approach [34] for extrapolating occupancy maps beyond visible regions (with a relative gain of 8% in floorplan accuracy). Though observing only a small fraction of the full homes, our model yields good interior maps covering much of their area. We also show audio and vision are synergistic signals to classify room types, allowing high-level perception of the semantics of the space even before directly visiting each room.

2. Related Work

Floorplan and room layout reconstruction The vision and graphics communities have explored various ways to use visual data, depth sensors, and laser scanners to build floorplans. Geometric approaches use 3D point cloud inputs to

construct building-wide floor plans [42, 31]. Given RGB-D scans, FloorNet [29] and Floor-SP [8] estimate 2D floorplans and rooms’ semantic labels using a mix of deep learning and optimization. Given monocular RGB images [27] or 360° panoramas [46, 49, 43, 50], other methods estimate a 3D indoor Manhattan room layout. Using only a small portion of a 360° panorama, models can be trained to infer missing viewpoints [23] and/or semantic labels [40]. Unlike any of the above, our approach leverages both audio and visual sensing to infer a 2D floorplan map and its semantic room labels. As our results show, audio offers the advantage of sensing further beyond the field of view of visual sensors.

Mapping for navigation With adequate overlapping views, structure-from-motion methods can recover the 3D structure of an environment (e.g., [2, 38]). Laser-based 2D SLAM is often used in mobile robotics to obtain the ground plane map [35]. Recent work leverages scans of indoor environments [3, 41] and fast simulation tools [37] to facilitate work on embodied visual navigation [20, 36, 25, 4, 9, 34]. Often methods estimate an explicit 2D occupancy map, projecting the observed point cloud to the ground plane and growing the map over time [9, 34, 4]. To navigate to a specified room, the method of [30] predicts a 2D semantic map with room labels, learning their layout patterns. In contrast to navigation, where an intelligent agent controls the camera and builds its map in service of reaching a target, our goal is to transform a passive video sequence (with audio) into a map. We show the advantages of our audio-visual approach over OccAnt [34], the state-of-the-art navigation method that extrapolates beyond visible scene points using vision alone.

Audio for spatial sensing Prior work explores ways to exploit audio alone to sense the shape of a room or object. Given multiple microphone recordings of a known sound, the method of [13] computes the shape of a single convex polyhedral room, while sound reflections are used to sense the 3D shape of an object hidden around a corner [28]. In robotics, echolocation can detect distances to surfaces based on the reflections [39, 14, 45, 10]. Multi-channel audio is also used to track dynamic objects [11, 16]. Unlike any of these methods, our approach takes a video (both the audio and visual streams) as input and produces a floorplan as output. Furthermore, our model is not restricted to known microphone layouts or known emitted sounds; rather, it can learn from natural sounds sensed passively in the environment (e.g., running water, door shutting). While environment semantics are explored in acoustic scene analysis [18, 19], our problem is quite different: the target output is a geometric map, not a label for the acoustic event that occurred.

Audio-visual spatial sensing Audio and vision together offer powerful cues for spatial perception. At the object level, they reveal shape and material properties [48, 32], e.g., via the sound of one object striking another. At the environment

level, audio can help sense 3D surfaces in cases where depth sensing would fail, e.g., transparent, shiny, or textureless surfaces [47, 26], provide self-supervisory cues [17], and refine visual estimates of room geometry [22]. Recent work also leverages audio-visual sensing to address navigation tasks [6, 15, 7, 12, 5]. None of the above methods produce audio-visual floorplans. Furthermore, an insight unique to our work is the use of naturally occurring semantic sounds to understand a multi-room layout.

3. Approach

Our goal is to estimate the 2D layout of an environment depicted in a short video. The 2D layout has two components: the structure of the interior area, and the semantic labels (room types) associated with each region of the interior. First, we formally describe the problem (Sec 3.1). We then describe our proposed model AV-Map (Sec 3.2) and introduce our training and inference procedure (Sec 3.3).

3.1. Problem Formulation

We consider videos generated by a camera and an ambisonic microphone following short trajectories through various home environments. Ambisonic audio is an order-limited physical characterization of a sound field using spherical harmonic basis functions [1, 33]. We represent a video by $V = \{(v_1, a_1), (v_2, a_2) \dots\}$ where v_t is the RGB frame and a_t is the audio clip sampled at time step t . We denote by $P_V = \{\mathbf{0}, r_2, r_3, \dots\}$ the position of the camera and microphone relative to the first time step in the coordinate system of the floorplan, where $r_i = (x_i, y_i, \theta_i)$ records the movement along the x- and y-axis on the 2D ground plane and θ_i is the rotation about the gravity axis. Relative pose changes in video can be estimated using computer vision [21]; for simplicity we assume camera poses are available for all methods. However, we find our model is robust to the range of odometry noise considered in the literature [34].

Each floorplan is parameterized as two variables: M_{int} and M_{room} , which represent the structure and semantics, respectively. The interior map M_{int} is a 2D binary grid that is a top-down view of the environment and represents the existence of floor, objects, and furniture by label 1, and walls and areas outside the environment by label 0. The room map M_{room} is a 2D grid taking $N_r + 1$ possible values with labels $\{1, \dots, N_r\}$ representing the N_r room types (kitchen, bathroom, etc.) and 0 representing walls and areas outside the environment. Each cell in the floorplan (an entry in the matrix M) represents a 25cm^2 area.

The goal of our work is to learn a mapping \mathcal{F} that estimates the floorplan (both M_{int} and M_{room}) of an environment using the video V and the relative pose changes P_V . The visual information in v_t captures the geometric properties and room types of visible regions. The audio information captured in a_t is either actively emitted by the camera, or

else passively generated by objects and people in the environment (details below). Since the placement of many objects is highly correlated with room types (for example, showers are in bathrooms and dishwashers are in kitchens), the audio signal can be a strong semantic signal indicating the room types. Furthermore, the echoes propagating through the environment capture geometric properties of the walls and other major surfaces. Our key insight is that the audio observations will illuminate the map for regions beyond what is visible in the frames of a short video.

3.2. AV-Map Floorplan Estimation Model

We now present our AV-Map floorplan estimation model \mathcal{F} . Fig. 2 overviews our proposed model, consisting of three components: Top-Down Feature Extraction, Feature Alignment, and a Sequence Encoder-Decoder architecture. At each time step, AV-Map estimates the interior map (M_{int}) and semantic room labels (M_{room}) in a neighborhood centered around the camera, integrating estimates over time.

Top-Down Feature Extraction The first stage of our proposed model involves extracting features for a given video V . The purpose of the feature extraction step is to project egocentric visual frames v_t and ambisonic audio clips a_t to a 2D feature grid that is spatially aligned with the top-down floorplan that we wish to estimate at each time step.

RGB Feature Extraction The feature extractor for an RGB frame v_t consists of a ResNet-18 model up to `layer2` followed by a spatial pooling operation, which leads to a single 128-D feature. We use `layer2` features of the ResNet to capture low-level features like corners of rooms. This feature is then upsampled by a sequence of transposed convolutions leading to the final visual feature $f_t \in \mathbb{R}^{128 \times \frac{H}{2} \times \frac{W}{2}}$ where H, W are the height and width of the considered output floorplan area at each time step. See Fig. 2, Column 1. Importantly, this predicted area extends beyond the freespace directly observable from the visual frame v_t ; in our experiments, the visually observed space on average covers only 14% of the area when $H \times W = 40\text{m}^2$ around the camera. Below we explain how maps from multiple steps are aggregated for the video’s final (larger) output map.

Audio Feature Extraction A video V also consists of audio clips $a_t \in \mathbb{R}^{T \times 9}$ where T is the sound duration at each time step t , and 9 is the number of ambisonic channels, corresponding to second order ambisonics. Features for each audio clip are extracted using a sequence of linear, ReLU, and pooling layers, yielding a 128-D feature. Similar to the RGB extractor, this feature is upsampled via transposed convolutions to obtain the final audio feature $g_t \in \mathbb{R}^{128 \times \frac{H}{2} \times \frac{W}{2}}$. See Fig. 2, Column 2.

Feature Alignment After extracting features for the RGB frames and audio clips, each video is represented by the set of visual features f_t and audio features g_t , for $t = 1, \dots, t_V$,

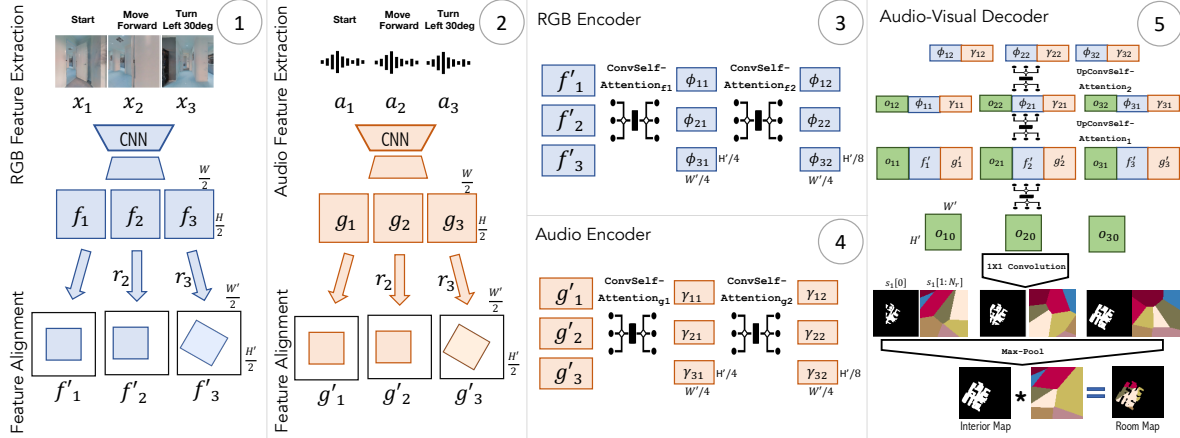


Figure 2: The AV-Map model has three stages: feature extraction, feature alignment, and sequence encoder-decoder. At each time step, the feature extractor obtains independent top-down floorplan-aligned features for each modality (ambisonic audio and RGB). These features are aligned to a common coordinate frame using the relative motion of the camera. Entire sequences of audio and visual features are then processed by an encoder using a series of pixelwise self-attention operations and convolution layers. Finally, the two modalities are fused using a decoder with a series of self-attention and convolution layers. We depict training for 3 time steps here. No matter how many steps are used in training, the model is applicable to arbitrary length sequences at test time.

where t_V is the number of frames in video V and may vary across videos. Note that each of these features was computed independently and thus far represents a feature aligned with the top-down map in a canonical *frame-centric* coordinate frame. In order to process entire sequences, we need to establish correspondences between the features at each time step. Therefore, we next align all the features to a common coordinate system, relative to the first frame. In order to retain relative pose information, we concatenate a 64-channel 2D positional encoding map [44] to each of the features (see Supp. for more details). The aligned visual and audio features $f'_t, g'_t \in \mathbb{R}^{(128+64) \times \frac{H'}{2} \times \frac{W'}{2}}$ are computed by padding with zeros, and translating and rotating each feature f_t, g_t by r_t where $H' > H, W' > W$ due to padding. See Fig. 2, bottom of Columns 1 and 2.

Sequence Encoder-Decoder We now wish to encode features for each time step that account for information present in the entire sequence. For example, the appearance of a wall in the second frame should inform the features in the first step and vice-versa. Self-attention [44] operations have shown to be useful to encode such bi-directional relationships. Inspired by this, we design a sequence of two self-attention and convolution operations (as shown in Fig. 2) which we refer to as *ConvSelfAttention*. The self-attention operations are responsible for communication across time steps at each pixel location (see Supp. for comparison to LSTM based model). We use convolutions with stride 2 to also simultaneously downsample the features. We denote the resulting features for the RGB frames as:

$$\phi_{t1} = \text{ConvSelfAttention}_{f1}(f'_t) \quad (1)$$

$$\phi_{t2} = \text{ConvSelfAttention}_{f2}(\phi_{t1}), \quad (2)$$

where ϕ_{t1}, ϕ_{t2} represent the intermediate encoder features. Similarly for the audio data, we represent by

γ_{t1}, γ_{t2} the outputs of the corresponding encoding layers *ConvSelfAttention*_{g1} and *ConvSelfAttention*_{g2}. Note that since the convolutions downsample the features, we have $\phi_{t1}, \gamma_{t1} \in \mathbb{R}^{C_1 \times \frac{H'}{4} \times \frac{W'}{4}}$ and $\phi_{t2}, \gamma_{t2} \in \mathbb{R}^{C_2 \times \frac{H'}{8} \times \frac{W'}{8}}$.

So far, we have processed the visual and audio information independently. In order to take full advantage of the presence of both modalities, we allow cross-modal information transfer. We accommodate this in the final decoding stage by concatenating the corresponding intermediate visual and audio features. For the decoder, we follow an architecture similar to the encoder by replacing the convolutions with transposed convolutions to upsample the features. We refer to these layers as *UpConvSelfAttention*. More concretely, the decoder consists of three *UpConvSelfAttention* layers which are used to compute the output as:

$$o_{t2} = \text{UpConvSelfAttention}_2([\phi_{t2}, \gamma_{t2}]) \quad (3)$$

$$o_{t1} = \text{UpConvSelfAttention}_1([o_{t2}, \phi_{t1}, \gamma_{t1}]) \quad (4)$$

$$o_{t0} = \text{UpConvSelfAttention}_0([o_{t1}, f'_t, g'_t]). \quad (5)$$

The final output $o_{t0} \in \mathbb{R}^{C \times H' \times W'}$ is classified using a 1x1 convolution, giving the final prediction at each time step as $s_t \in \mathbb{R}^{(N_r+1) \times H' \times W'}$. The first channel (notated as $s_t[0, :, :]$) represents a binary score map for the existence of interior space, and the remaining channels 1, 2, ..., N_r (notated as $s_t[i, :, :]$) represent score maps for the existence of the corresponding room type. Note that due to the alignment step above, these output maps are aligned in the common coordinate frame of the first time step. Therefore, to produce a prediction S for the whole sequence, we max-pool the predictions s_t . The self-attention in the earlier encoder-decoder already accounts for communication across time steps for these per-step estimates. AV-Map outputs the aggregated interior and room classification scores for a video sequence:

$$S = \mathcal{F}(v_1, a_1, v_2, a_2, \dots, v_{t_V}, a_{t_V}).$$

In practice, for training, we fix the length of sequences $t_V = 4$, which balances memory constraints with learning to integrate over time. For illustration, Fig. 2 depicts an instance of the model with $t_V = 3$.

In summary, the proposed AV-Map floorplan estimation model processes audio-visual sequences at various levels. The feature extraction independently processes each time step. The top-down alignment brings the features to a common coordinate frame. The encoders process sequences of each modality independently while integrating information across time, and finally the decoder fuses information from both visual and audio modalities.

3.3. Training and Inference

The output of AV-Map is a 2D map with $N_r + 1$ channels. The model is trained to predict two floorplan maps: the interior structure and the pixel-wise room labels.

Predicting Interior Maps Prediction of interior maps is a pixel-wise binary classification problem where 0s represent the walls or exterior points and 1s represent the points inside the environment (floors, furniture, objects, etc.). From s_t , the pixel-wise binary classification probability is computed using the sigmoid function: $p_{t_{\text{int}}}[i, j] = \frac{1}{1 + \exp -s_t[0, i, j]}$ for each pixel location i, j in the 2D grid.

Predicting Room Floorplans Prediction of room floorplans is similar to prediction of interior maps, but requires multi-class classification of each pixel into one of N_r semantic room types. Therefore, the class-wise probabilities at each pixel are computed using the softmax function. Concretely, the classification probability for class $n \in \{1, 2, 3, \dots, N_r\}$ at pixel location (i, j) is computed as: $p_{t_{\text{room}}}[n - 1, i, j] = \frac{\exp s_t[n, i, j]}{\sum_{k=1}^{N_r} \exp s_t[k, i, j]}$.

Training Objectives For each time step t , let the ground truth interior and room maps of the $H \times W$ area around the camera be represented by $y_{t_{\text{int}}} \in \{0, 1\}^{H \times W}$ and $y_{t_{\text{room}}} \in \{0, 1, \dots, N_r\}^{H \times W}$. Since our model’s predictions are aligned with time step $t = 1$, we similarly align the ground truth maps to obtain $\hat{y}_{t_{\text{int}}} \in \{0, 1\}^{H' \times W'}$ and $\hat{y}_{t_{\text{room}}} \in \{0, 1, \dots, N_r\}^{H' \times W'}$ by padding with zeros, translating and rotating by r_t (where H', W' are the increased dimensions due to padding). The interior and room map classification objectives for each time step t for pixel location (i, j) are then defined as:

$$\mathcal{L}_{\text{int}} = \frac{1}{z} \sum_{t=1}^{t_V} \sum_{i=1}^{H'} \sum_{j=1}^{W'} -\hat{y}_{t_{\text{int}}}[i, j] \log p_{t_{\text{int}}}[i, j] \quad (6)$$

$$\mathcal{L}_{\text{room}} = \frac{1}{z} \sum_{t=1}^{t_V} \sum_{i=1}^{H'} \sum_{j=1}^{W'} \sum_{k=1}^{N_r} -\mathbb{I}[\hat{y}_{t_{\text{room}}}[i, j]=k] \log p_{t_{\text{room}}}[k, i, j], \quad (7)$$

where $z = t_V H' W'$, \mathbb{I} is the indicator function, and t_V is the number of time steps in the video V . We ignore the unused pixel locations (i, j) in \hat{y}_t that arise from padding during the alignment step. AV-Map is trained using the sum of the two objectives: $\mathcal{L} = \mathcal{L}_{\text{int}} + \mathcal{L}_{\text{room}}$.

During inference, we estimate the interior and room maps for the whole sequence. As explained above, this is done by max-pooling the predictions s_t to produce a sequence-level prediction S . Importantly, the self-attention layers in our model ensure that entire sequences are used to reason about each time step. Furthermore, since self-attention layers can process sequences of arbitrary length, we can apply the trained model on videos of varying length.

In order to predict the binary interior map, we simply threshold at $p = 0.5$ the final pixel-wise interior probabilities. To obtain the room map prediction, we assign the most likely room label to each location and use the thresholded interior map prediction as a binary mask to get its shape.

3.4. Video Sequence Generation

In order to generate videos in a variety of 3D environments for which we know ground truth floorplans, we use the Matterport3D dataset [3]¹ and the SoundSpaces [6] audio simulations. SoundSpaces provides highly realistic audio for 85 fully scanned real environments split 59/11/15 for train/val/test, respectively. Most environments are large multi-room homes and contain a variety of furnishings. SoundSpaces provides precomputed impulse responses (IR) for all source-receiver locations on a dense grid sampled at 1m spatial resolution. The simulations use SoTA multi-band ray tracing, computing the IRs from arbitrary geometries and frequency-dependent acoustic material properties, and modeling both transmission (including through walls) and scattering. The IRs can be convolved with any audio clip to generate realistic audio for any chosen source-receiver location, including multiple simultaneous sources. See [6] for details of the simulations and Supp. videos for examples.

Generating Floorplans We use the Habitat-API [37] to generate top-down interior floorplans for each environment by projecting the point cloud to the 2D ground plane. Room floorplans are constructed using the Matterport3D room annotations by assigning a room label to each pixel of the interior floorplan. We use the $N_r=13$ most frequent room labels from Matterport3D (laundry, kitchen, bathroom, etc.).

Camera Trajectories We generate videos by recording ego-centric frames and ambisonic audio along short camera trajectories. Due to the grid constraint of the SoundSpaces IRs, we restrict camera positions to the same 1m grid. At each location, the camera is parallel to the ground plane and can have a rotation around the gravity-axis in the set $\{0^\circ, 30^\circ, 60^\circ, \dots, 330^\circ\}$. The trajectories are generated by

¹Matterport3D license: http://kaldir.vc.in.tum.de/matterport/MP_TOS.pdf.

randomly sampling actions from { ‘forward’, ‘left 30°’, ‘right 30°’ } with probabilities {0.8, 0.1, 0.1}, respectively, to avoid unnaturally frequent loops or oscillations. At each step, with the RGB frame constant, we record audio for 3 seconds.

We train the models with fixed trajectory lengths of $t_V = 4$ steps due to GPU memory constraints. In Supp. we provide an ablation with $t_V = 1$ to demonstrate the power of learning across the sequence. For evaluation on the validation and test environments, we sample trajectories of variable length t_V , for $t_V \in \{1, 2, 4, 8, 16\}$.

Audio We consider two settings of audio: *device-generated* (active) and *environment-generated* (passive). For the device-generated (Dev. Gen.) setting, the video recording device (e.g., cell phone, AR headset, or robot) also emits a fixed recurring sound at each time step. We use a 3 sec frequency sweep chirp signal in the audible range (20Hz-20KHz). Though any emitted sound could provide useful echoes, the wide range of frequencies activated in the sweep is expected to provide a particularly rich learning signal [24].

In the environment-generated setting, rather than emit a sound, the system listens for naturally occurring sounds in homes. To achieve this, we first collect a set of 56/32/32 train/val/test audio clips² of duration 3 sec that capture sounds made by objects in different room types (for example, sound of a flush, dishwasher, etc.). This allows us to place source sounds in the Matterport3D environments in the appropriate rooms. For each trajectory, the location of sound source(s) is randomly chosen, and the waveform played is dependent on the room type of that location.

We consider three passive settings: 1) Env. Telephone: the source is near (within 40m² area) one of the steps in the trajectory and plays the “telephone” sound; 2) Env. Nearby: there is again a single sound source near the trajectory, but the audio clip varies according to the room type of the sampled source location; 3) Env. All Room: a source is randomly placed in each room and all sources simultaneously play a sound associated with their room type.

4. Results

Through extensive qualitative and quantitative results, we demonstrate that our proposed model can effectively leverage both audio and visual signals to reason about the extent of the interior of environments (Sec 4.1) and classify regions of the interior into the associated rooms (Sec 4.2).

Baselines In order to conduct a thorough analysis of our proposed model, we consider several baselines.

Interior-only: simple baseline that predicts interior pixels (1s) everywhere in the considered neighborhood.

Projected Depth: standard occupancy map computed by projecting depth maps to the ground plane [9, 4]. Note that our AV model does not leverage depth, only RGB and audio.

²Downloaded from freesound.org

Table 1: Interior reconstruction: Our AV-Map model (here with device-generated sounds) outperforms existing methods and the baselines. Methods creating only a binary map output cannot be scored by AP (NA). Standard errors are maximally 0.0018; see Supp. for details.

	AP	Acc.	Edge AP
Interior-only	NA	50.00	NA
Projected Depth	NA	53.73	NA
OccAnt [34]	60.27	58.45	51.52
OccAnt [34] +40m ² Area	68.24	62.89	51.88
Acoustic Echoes [13]	NA	50.37	NA
(Ours) RGB-only	71.07	65.15	53.90
(Ours) RGB+D	71.70	65.60	53.34
(Ours) Audio-only	69.99	64.83	53.20
(Ours) Audio-Visual	73.28	66.52	54.67

OccAnt [34]: The SoTA Occupancy Anticipation model [34] infers an interior map (at each time step) from RGB-D by learning to extrapolate beyond the visible ground-plane projections. It is a key baseline to test our claim that audio can better “see” beyond the visual observations. We use the authors’ code, since their reported results test only on Gibson, for which sound simulations do not exist. For fair comparison, we report results both using OccAnt exactly as it was proposed (estimating 9m² in front of the camera [34]) as well as retraining OccAnt to estimate 40m² around the camera (same area as our model).

Acoustic Echoes [13]: This method assumes that all room shapes are convex polyhedra and estimates room shape by listening to audio echoes. However, this approach requires knowing the *ground-truth* impulse responses at each microphone location, which our method does *not* have access to. While this method’s setup is artificial, we use it as an upper bound for what an existing audio-only method can provide.

Ours audio-only and RGB-only: As ablations of our model, we train variants with either modality removed.

Note that existing models like FloorNet [29, 8] are not applicable, since they require fully scanned point clouds as input. In our setting, the input is simply a short sequence of egocentric RGB views and audio.

4.1. Floorplan Interior Reconstruction

First we present interior floorplan results. We set $H \times W$ such that it covers 40m² area at each time step (see Supp. for similar results with 164m²). Since we aggregate the predictions from all time steps, the final accumulated area varies with the number of steps and direction of movement, adding at most 6.25m² at each step, for final output areas ranging from 40m² (1 step) to 134m² (16 steps).

Evaluation Metrics We use three metrics: Average Precision (AP), Accuracy (Acc.), and Edge Average Precision (Edge AP). AP and Acc compare $S[0, :, :]$ and the binary ground truth map. Edge AP compares the edges of the predicted and ground truth maps in order to emphasize differences in boundary shapes. Pixels are reweighted in all

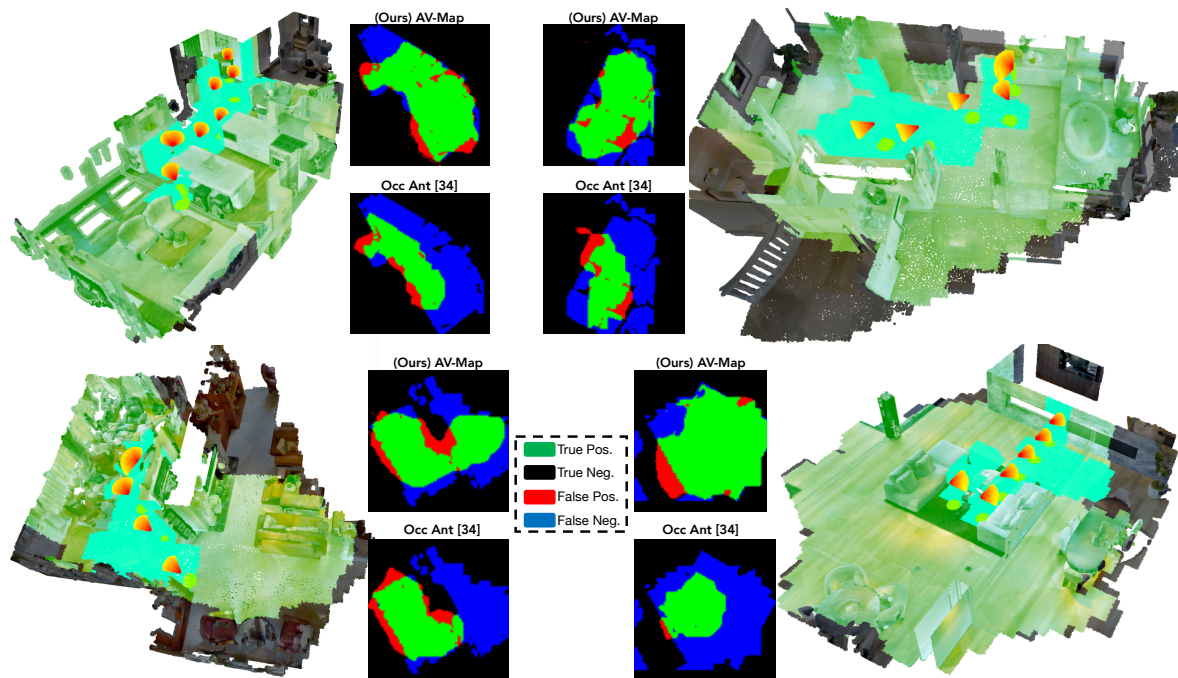


Figure 3: Interior map predictions: Here we visualize reconstructed interior maps estimated by OccAnt [34] and our proposed audio-visual model in the device-generated audio setting. In the 3D examples, red cones are camera positions, cyan is the observed ground plane area, and transparent light green is the freespace inferred by our model. 2D output maps are color-coded against ground truth (see legend). The perfect 2D map would be all green and black. Our method “sees” substantially more area by using audio, and it produces more accurate maps than the state-of-the-art mapping method [34] that also attempts to extrapolate beyond the directly observed area.




Figure 4: AV-Map vs. RGB-only: Example interior map predictions of AV-Map and its RGB-only variant. In red, we show interior locations correctly captured by AV-Map but not captured by the RGB-only model. In blue, we show locations captured by the RGB-only model but missed by AV-Map. Observe that AV-Map anticipates regions that are not visible through the camera (cone), which the RGB-only model cannot capture.

metrics to balance the contribution of labels 0 and 1.

Comparison to Baselines Tab. 1 presents our central result, a quantitative evaluation of all models on test trajectories of length 4 steps in unseen environments. Our AV-Map model outperforms all the baselines on all three metrics. Furthermore, our full AV model outperforms the audio-only and RGB(+/-D)-only variants by good margins. This shows our model successfully performs joint inference by leveraging important cues from both modalities. Our model with RGB-only/Audio-only is itself stronger than existing methods for either modality [34, 13], showing the strength of our proposed framework even without the advantage of more modalities. Fig. 3 shows example map predictions compared to [34], the best existing method. They highlight how audio allows “seeing” both behind the camera as well as inferring freespace behind walls in large multi-room homes. Fig. 4 compares AV-Map and its RGB-only variant.

Comparison of Audio Settings Table 2 compares our AV-Map model in the three audio settings described in Sec 3.4. The device-generated allows us to capture structure near the camera and the environment-generated audio captures information about adjacent or far-away rooms. Hence both offer complementary strengths. Our model performs better when operated in the device-generated audio setting. This is also expected since the frequency sweep audio allows us to capture all the frequencies in the audible range, unlike the naturally occurring sounds in the




Figure 5: Effect of trajectory length

Table 2: Interior reconstruction in different settings of audio: AV-Map is applicable with either device-generated or environment-generated sounds.

		AP	Acc.	Edge AP
	RGB-only	71.07	65.15	53.90
Dev. Gen.	Audio-only	69.99	64.83	53.20
	Audio-Visual	73.28	66.52	54.67
Env. Telephone	Audio-only	68.60	63.53	53.27
	Audio-Visual	72.30	66.69	54.16
Env. Nearby	Audio-only	66.27	61.49	51.86
	Audio-Visual	72.86	66.16	54.32
Env. All Room	Audio-only	65.42	61.41	51.98
	Audio-Visual	71.38	65.46	54.79

environment-generated settings. Furthermore, the relative location of the source is known in the device-generated setting (since it is always at the camera). In every setting, AV-Map outperforms the RGB-only and Audio-only ablations. Despite the challenges in Env. All Room (multiple simultaneous sounds from different source locations), we observe minimal decline in our interior map performance.

Effect of Trajectory Length Fig. 5 shows our *full-house* accuracy gains as a function of trajectory length.³ Our model consistently outperforms the RGB-only ablation consistently. Notably, our gains are largest when the video is shorter, so less area is visible. This again confirms the power of audio to “see” beyond the images. As the length increases, the visible fraction of the environment becomes larger, diminishing the impact of audio signals (dotted blue line). See Supp. for ablation showing the impact of our model’s self-attention.

4.2. Floorplan Room Classification

Finally, we evaluate the semantic room label maps. We evaluate with Env-Gen. audio, which provides natural object-room cues. Fig. 6 compares the mean of the pixel-wise room classification AP for our model in the Env-Gen. All Room setting and its ablations.⁴ Results are averaged over

³Absolute numbers are lower than in Tab. 1 because the scored area here is the entire house’s area, vs. the maximum output map area in Tab. 1.

⁴Note that the baselines from Table 1 are not applicable here because they produce only geometric interior maps.

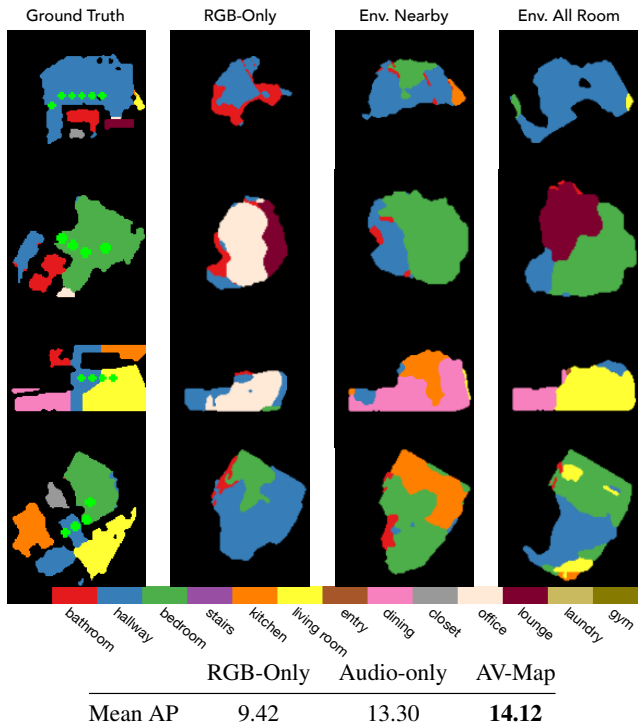


Figure 6: Room map predictions: AV-Map infers the layout of rooms in the home based on what it sees in the short video as well as the natural occurring sounds it hears. Green dots on GT maps are camera positions.

all trajectory lengths. Audio helps identify rooms around the trajectory, and we see best results with both modalities. The room map examples show AV-Map provides better room classification compared to RGB-only, and does best with many natural semantic sounds—encouraging for deployment in a busy household. Our model can identify the correct room type and its approximate location, though without actually entering a room, its exact footprint remains ambiguous.

5. Conclusion and Future Work

We proposed a new research direction: audio-visual floorplan reconstruction from short video sequences. We developed a multi-modal model to estimate the floorplan around and far beyond the camera trajectory. Our AV-Map model successfully infers the structure and semantics of areas that are not visible, outperforming the state-of-the-art in extrapolated visual maps. In future work we plan to consider extensions to multi-level floorplans and connect our mapping idea to a robotic agent actively controlling the camera.

6. Acknowledgements

We would like to thank Unnat Jain, Changan Chen, and Santhosh Kumar Ramakrishnan for help with generating the audio simulations and providing help with code for this work. We would also like to thank Ruohan Gao for providing feedback on the text.

References

- [1] T. Abhayapala and D. Ward. Theory and design of higher order sound field microphones using spherical microphone array. In *ICASSP*, 2002. 3
- [2] Marcus A. Brubaker, Andreas Geiger, and Raquel Urtasun. Lost! leveraging the crowd for probabilistic visual self-localization. In *CVPR*, 2013. 2
- [3] Angel Chang, Angela Dai, Thomas Funkhouser, Maciej Halber, Matthias Niessner, Manolis Savva, Shuran Song, Andy Zeng, and Yinda Zhang. Matterport3d: Learning from rgb-d data in indoor environments. *arXiv preprint arXiv:1709.06158*, 2017. 2, 5
- [4] Devendra Singh Chaplot, Dhiraj Gandhi, Saurabh Gupta, Abhinav Gupta, and Ruslan Salakhutdinov. Learning to explore using active neural slam. *arXiv preprint arXiv:2004.05155*, 2020. 2, 6
- [5] Changan Chen, Ziad Al-Halah, and Kristen Grauman. Semantic audio-visual navigation. In *CVPR*, 2021. 3
- [6] Changan Chen, Unnat Jain, Carl Schissler, Sebastia Vincenc Amengual Gari, Ziad Al-Halah, Vamsi Krishna Ithapu, Philip Robinson, and Kristen Grauman. SoundSpaces: Audio-visual navigation in 3d environments. In *ECCV*, 2020. 3, 5
- [7] Changan Chen, Sagnik Majumder, Ziad Al-Halah, Ruohan Gao, Santhosh Kumar Ramakrishnan, and Kristen Grauman. Learning to set waypoints for audio-visual navigation. In *ICLR*, 2021. 3
- [8] Jiacheng Chen, Chen Liu, Jiaye Wu, and Yasutaka Furukawa. Floor-sp: Inverse cad for floorplans by sequential room-wise shortest path. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2661–2670, 2019. 1, 2, 6
- [9] Tao Chen, Saurabh Gupta, and Abhinav Gupta. Learning exploration policies for navigation. In *7th International Conference on Learning Representations, ICLR 2019*, 2019. 2, 6
- [10] Jesper Haahr Christensen, Sascha Hornauer, and X Yu Stella. Batvision: Learning to see 3d spatial layout with two ears. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1581–1587. IEEE, 2020. 2
- [11] Marco Crocco, Samuele Martelli, Andrea Trucco, Andrea Zunino, and Vittorio Murino. Audio tracking in noisy environments by acoustic map and spectral signature. *EEE Trans Cybern*, 2018. 2
- [12] Victoria Dean, Shubham Tulsiani, and Abhinav Gupta. See, hear, explore: Curiosity via audio-visual association. In *NeurIPS*, 2020. 3
- [13] Ivan Dokmanić, Reza Parhizkar, Andreas Walther, Yue M Lu, and Martin Vetterli. Acoustic echoes reveal room shape. *Proceedings of the National Academy of Sciences*, 110(30):12186–12191, 2013. 2, 6, 7
- [14] Itamar Eliakim, Zahi Cohen, Gabor Kosa, and Yossi Yovel. A fully autonomous terrestrial bat-like acoustic robot. *PLoS computational biology*, 14(9):e1006406, 2018. 2
- [15] Chuang Gan, Yiwei Zhang, Jiajun Wu, Boqing Gong, and Joshua B Tenenbaum. Look, listen, and act: Towards audio-visual embodied navigation. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 9701–9707. IEEE, 2020. 3
- [16] Chuang Gan, Hang Zhao, Peihao Chen, David Cox, and Antonio Torralba. Self-supervised moving vehicle tracking with stereo sound. In *ICCV*, 2019. 2
- [17] R. Gao, C. Chen, Z. Al-Halah, C. Schissler, and K. Grauman. VisualEchoes: Spatial image representation learning through echolocation. In *ECCV*, 2020. 3
- [18] Jort Gemmeke, Daniel Ellis, Dylan Freedman, Aren Jansen, Wade Lawrence, Channing Moore, Manoj Plakal, and Marvin Ritter. Audio set: An ontology and human-labeled dataset for audio events. In *ICASSP*, 2017. 2
- [19] Dimitrios Giannoulis, Emmanouil Benetos, Dan Stowell, Mathias Rossignol, Mathieu Lagrange, and Mark D. Plumbley. Detection and classification of acoustic scenes and events: An ieeeaasp challenge. In *IEEE Workshop on Applications of Signal Processing to Audio and Acoustics*, 2013. 2
- [20] Saurabh Gupta, David Fouhey, Sergey Levine, and Jitendra Malik. Unifying map and landmark based representations for visual navigation. *arXiv preprint arXiv:1712.08125*, 2017. 2
- [21] R. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. 2004. 3
- [22] Muhammad Wajahat Hussain, Javier Civera, and Luis Montano. Grounding acoustic echoes in single view geometry estimation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 28, 2014. 3
- [23] D. Jayaraman and K. Grauman. Learning to look around: Intelligently exploring unseen environments for unknown tasks. In *CVPR*, 2018. 2
- [24] H. Jeong and Y. Lam. Source implementation to eliminate low-frequency artifacts in finite difference time domain room acoustic simulation. *Journal of the Acoustical Society of America*, 2012. 6
- [25] Abhishek Kadian, Joanne Truong, Aaron Gokaslan, Alexander Clegg, Erik Wijmans, Stefan Lee, Manolis Savva, Sonia Chernova, and Dhruv Batra. Are we making real progress in simulated environments? measuring the sim2real gap in embodied visual navigation. *arXiv preprint arXiv:1912.06321*, 2019. 2
- [26] Hansung Kim, Luca Remaggi, Philip JB Jackson, Filippo Maria Fazi, and Adrian Hilton. 3d room geometry reconstruction using audio-visual sensors. In *2017 International Conference on 3D Vision (3DV)*, pages 621–629. IEEE, 2017. 3
- [27] Chen-Yu Lee, Vijay Badrinarayanan, Tomasz Malisiewicz, and Andrew Rabinovich. Roomnet: End-to-end room layout estimation. In *ICCV*, 2017. 1, 2
- [28] David B Lindell, Gordon Wetzstein, and Vladlen Koltun. Acoustic non-line-of-sight imaging. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6780–6789, 2019. 2
- [29] Chen Liu, Jiaye Wu, and Yasutaka Furukawa. Floornet: A unified framework for floorplan reconstruction from 3d scans. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 201–217, 2018. 1, 2, 6
- [30] Medhini Narasimhan, Erik Wijmans, Xinlei Chen, Trevor Darrell, Dhruv Batra, Devi Parikh, and Amanpreet Singh.

- Seeing the un-scene: Learning amodal semantic maps for room navigation. *arXiv preprint arXiv:2007.09841*, 2020. 2
- [31] Brian Okorn, Xuehan Xiong, Burcu Akinci, and Daniel Huber. Toward automated modeling of floor plans. In *3DPVT*, 2009. 1, 2
- [32] Andrew Owens, Phillip Isola, Josh McDermott, Antonio Torralba, Edward H Adelson, and William T Freeman. Visually indicated sounds. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2405–2413, 2016. 2
- [33] B. Rafaely. Analysis and design of higher order sound field microphones using spherical microphone array. *IEEE Trans. on Speech and Audio Processing*, 2005. 3
- [34] Santhosh K Ramakrishnan, Ziad Al-Halah, and Kristen Grauman. Occupancy anticipation for efficient exploration and navigation. *arXiv preprint arXiv:2008.09285*, 2020. 2, 3, 6, 7
- [35] J. Santos, D. Portugal, and R. Rocha. An evaluation of 2d slam techniques available in robot operating system. In *IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR)*, 2013. 2
- [36] Nikolay Savinov, Alexey Dosovitskiy, and Vladlen Koltun. Semi-parametric topological memory for navigation. *arXiv preprint arXiv:1803.00653*, 2018. 2
- [37] Manolis Savva, Abhishek Kadian, Oleksandr Maksymets, Yili Zhao, Erik Wijmans, Bhavana Jain, Julian Straub, Jia Liu, Vladlen Koltun, Jitendra Malik, et al. Habitat: A platform for embodied ai research. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 9339–9347, 2019. 2, 5
- [38] S. Se, D. Lowe, and J. Little. Mobile robot localization and mapping with uncertainty using scale-invariant visual landmarks. *IJRR*, 2002. 2
- [39] Jascha Sohl-Dickstein, Santani Teng, Benjamin M Gaub, Chris C Rodgers, Crystal Li, Michael R DeWeese, and Nicol S Harper. A device for human ultrasonic echolocation. *IEEE Transactions on Biomedical Engineering*, 62(6):1526–1534, 2015. 2
- [40] Shuran Song, Andy Zeng, Angel X. Chang, Manolis Savva, Silvio Savarese, and Thomas Funkhouser. Im2pano3d: Extrapolating 360 structure and semantics beyond the field of view. In *CVPR*, 2018. 2
- [41] Julian Straub, Thomas Whelan, Lingni Ma, Yufan Chen, Erik Wijmans, Simon Green, Jakob J Engel, Raul Mur-Artal, Carl Ren, Shobhit Verma, et al. The replica dataset: A digital replica of indoor spaces. *arXiv preprint arXiv:1906.05797*, 2019. 2
- [42] Wei Sui, Lingfeng Wang, Bin Fan, Hongfei Xiao, Huaiyu Wu, and Chunhong Pan. Layer-wise floorplan extraction for automatic urban building reconstruction. *IEEE Transactions on Visualization and Computer Graphics*, 2016. 1, 2
- [43] Cheng Sun, Chi-Wei Hsiao, Min Sun, and Hwann-Tzong Chen. Horizonnet: Learning room layout with 1d representation and pano stretch data augmentation. In *CVPR*, 2019. 2
- [44] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017. 4
- [45] Antonio Pico Villalpando, Guido Schillaci, Verena V Hafner, and Bruno Lara Guzmán. Ego-noise predictions for echolocation in wheeled robots. In *Artificial Life Conference Proceedings*, pages 567–573. MIT Press, 2019. 2
- [46] S. Yang, F. Wang, C. Peng, P. Wonka, M. Sun, and H. Chu. Dula-net: A dual-projection network for estimating room layouts from a single rgb panorama. In *CVPR*, 2019. 1, 2
- [47] Mao Ye, Yu Zhang, Ruigang Yang, and Dinesh Manocha. 3d reconstruction in the presence of glasses by acoustic and stereo fusion. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4885–4893, 2015. 3
- [48] Zhoutong Zhang, Jiajun Wu, Qiuqia Li, Zhengjia Huang, James Traer, Josh H McDermott, Joshua B Tenenbaum, and William T Freeman. Generative modeling of audible shapes for object perception. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1251–1260, 2017. 2
- [49] Chuhan Zou, Alex Colburn, Qi Shan, and Derek Hoiem. Layoutnet: Reconstructing the 3d room layout from a single rgb image. In *CVPR*, 2018. 2
- [50] Chuhan Zou, Jheng-Wei Su, Chi-Han Peng, Alex Colburn, Qi Shan, Peter Wonka, Hung-Kuo Chu, and Derek Hoiem. 3d manhattan room layout reconstruction from a single 360 image. *arXiv preprint arXiv:1910.04099*, 2019. 2